

NASA-CR-193688

JOHNSON

10-82-CN

186043

39 p

NASA Electronic Library System (NELS) Optimization

William L. Pribyl
DataCraft, Inc.

May 1992 - July 1993

N94-18839


Unclas

G3/82 0186043

Cooperative Agreement NCC 9-16
Research Activity No. RB.07:
RBSE: ASV3/NELS Database Optimization Program

NASA Technology Utilization Program
NASA Headquarters

(NASA-CR-193688) NASA ELECTRONIC
LIBRARY SYSTEM (NELS) OPTIMIZATION
(Houston Univ.) 39 p



Research Institute for Computing and Information Systems
University of Houston-Clear Lake

DELIVERABLE REPORTS

The RICIS Concept

The University of Houston-Clear Lake established the Research Institute for Computing and Information Systems (RICIS) in 1986 to encourage the NASA Johnson Space Center (JSC) and local industry to actively support research in the computing and information sciences. As part of this endeavor, UHCL proposed a partnership with JSC to jointly define and manage an integrated program of research in advanced data processing technology needed for JSC's main missions, including administrative, engineering and science responsibilities. JSC agreed and entered into a continuing cooperative agreement with UHCL beginning in May 1986, to jointly plan and execute such research through RICIS. Additionally, under Cooperative Agreement NCC 9-16, computing and educational facilities are shared by the two institutions to conduct the research.

The UHCL/RICIS mission is to conduct, coordinate, and disseminate research and professional level education in computing and information systems to serve the needs of the government, industry, community and academia. RICIS combines resources of UHCL and its gateway affiliates to research and develop materials, prototypes and publications on topics of mutual interest to its sponsors and researchers. Within UHCL, the mission is being implemented through interdisciplinary involvement of faculty and students from each of the four schools: Business and Public Administration, Education, Human Sciences and Humanities, and Natural and Applied Sciences. RICIS also collaborates with industry in a companion program. This program is focused on serving the research and advanced development needs of industry.

Moreover, UHCL established relationships with other universities and research organizations, having common research interests, to provide additional sources of expertise to conduct needed research. For example, UHCL has entered into a special partnership with Texas A&M University to help oversee RICIS research and education programs, while other research organizations are involved via the "gateway" concept.

A major role of RICIS then is to find the best match of sponsors, researchers and research objectives to advance knowledge in the computing and information sciences. RICIS, working jointly with its sponsors, advises on research needs, recommends principals for conducting the research, provides technical and administrative support to coordinate the research and integrates technical results into the goals of UHCL, NASA/JSC and industry.

***NASA Electronic Library System
(NELS) Optimization***



RICIS Preface

This research was conducted under auspices of the Research Institute for Computing and Information Systems by William L. Pribyl of DataCraft, Inc. Dr. E. T. Dickerson served as RICIS research coordinator.

Funding was provided by the NASA Technology Utilization Program, NASA Headquarters, Code C, through Cooperative Agreement NCC 9-16 between the NASA Johnson Space Center and the University of Houston-Clear Lake. The NASA research coordinator for this activity was Ernest M. Fridge III, Deputy Chief of the Software Technology Branch, Information Technology Division, Information Systems Directorate, NASA/JSC.

The views and conclusions contained in this report are those of the author and should not be interpreted as representative of the official policies, either express or implied, of UHCL, RICIS, NASA or the United States Government.



TABLE OF CONTENTS

- NELS Optimization: May 92 Progress/Problem Report**
DataCraft, Inc., 5/29/92
- NELS Optimization: June 92 Incremental Progress/Problem Report**
DataCraft, Inc., 6/30/92
- NELS Optimization: July 92 Incremental Progress/Problem Report**
DataCraft, Inc., 8/9/92
- NELS Optimization: Final Report, Phase 2 Tasks**
DataCraft, Inc., 8/30/92
- NELS Optimization: Interim Report, Phase 3**
DataCraft, Inc., 9/30/92
- NELS Optimization: Final Report, Phase 3**
DataCraft, Inc., 10/31/92
- NELS Optimization: Interim Report, Phase 4**
DataCraft, Inc., 11/30/92
- NELS Optimization: Interim Report Phase 4**
DataCraft, Inc., 1/11/93
- NELS Optimization: Progress/Problem, Phase 4**
DataCraft, Inc., 2/7/93
- NELS Optimization: Final Report, Phase 4**
DataCraft, Inc., 3/9/93
- NELS Optimization: Post Final Report, Phase 4**
DataCraft, Inc., 6/3/93
- NELS Optimization: Final, Last Phase**
DataCraft, Inc., 7/12/93

1. The first part of the document is a letter from the President of the United States to the Congress, dated January 3, 1862. It is a very important document, as it contains the President's annual message to Congress. The letter is written in a formal, dignified style, and it is one of the most important documents in the history of the United States.

2. The second part of the document is a letter from the Secretary of the Treasury to the Congress, dated January 3, 1862. It is a very important document, as it contains the Secretary's report on the state of the Treasury. The letter is written in a formal, dignified style, and it is one of the most important documents in the history of the United States.

3. The third part of the document is a letter from the Secretary of the Interior to the Congress, dated January 3, 1862. It is a very important document, as it contains the Secretary's report on the state of the Interior. The letter is written in a formal, dignified style, and it is one of the most important documents in the history of the United States.

4. The fourth part of the document is a letter from the Secretary of the War to the Congress, dated January 3, 1862. It is a very important document, as it contains the Secretary's report on the state of the War. The letter is written in a formal, dignified style, and it is one of the most important documents in the history of the United States.

5. The fifth part of the document is a letter from the Secretary of the Navy to the Congress, dated January 3, 1862. It is a very important document, as it contains the Secretary's report on the state of the Navy. The letter is written in a formal, dignified style, and it is one of the most important documents in the history of the United States.

6. The sixth part of the document is a letter from the Secretary of the State to the Congress, dated January 3, 1862. It is a very important document, as it contains the Secretary's report on the state of the State. The letter is written in a formal, dignified style, and it is one of the most important documents in the history of the United States.

7. The seventh part of the document is a letter from the Secretary of the War to the Congress, dated January 3, 1862. It is a very important document, as it contains the Secretary's report on the state of the War. The letter is written in a formal, dignified style, and it is one of the most important documents in the history of the United States.

8. The eighth part of the document is a letter from the Secretary of the Navy to the Congress, dated January 3, 1862. It is a very important document, as it contains the Secretary's report on the state of the Navy. The letter is written in a formal, dignified style, and it is one of the most important documents in the history of the United States.

NELS Optimization: May 92 Progress/Problem Report

Issues (in order of importance)

1. It is critical to obtain access to key NELS personnel to walk through the design and phrasing of those SQL statements that we have found to be bottlenecks. Specifically, we request a minimum two-hour meeting with C. J. Millebeck and Mark Rorvig to obtain authoritative decisions on whether certain changes fall within the bounds of NELS requirements. Specific areas of concern include rephrasing SQL statements generated by boolean queries and reorganizing "expensive" queries into a two-stage approach.
2. It is important to obtain and install the target hardware by mid-June in order that it can be configured and tuned by the end of the month. System level tuning must be performed *in situ* or is likely to be compromised.
3. Access to the lab has been a problem. On one weekday we were told that the lab was closed to us. On another occasion we were told that weekend work was not possible because of locked doors. This runs counter to the instructions we received that the vast majority of the work must be performed on site. Additionally, when tuning performance issues it is important to be able to have exclusive use of the system in order to isolate problems.
4. We have been told to work on our own copy of the Pro*C code, but a mechanism has not yet been established by which our code changes will be rolled into the production version.
5. Access to workstations has been another problem. We find that we generally cannot use the same machine on two consecutive visits to the lab, which consumes unnecessary time in learning a given workstation's idiosyncrasies.

Progress on Items Listed in Statement of Work

Each item from the statement of work is listed in italics, with a progress report below it.

1. *Examine NELS database, particularly in the memory, disk contention and CPU to discover bottlenecks.*

DataCraft has obtained performance requirements from Karen Fleming/Mountain Net and begun informal benchmarks with a set of test data loaded at our request by other NELS development personnel (Wes White et al). Initial results indicated several areas where performance appears to fall short of what is desired. For example, natural language queries seem to resolve in approximately ten seconds rather than the five seconds (five is the target for keyword search). Boolean queries take considerably longer than desired.

We have run a series of performance monitoring tools including vmstat, top, bstat/estat, ps, tkprof, and Explain Plan to isolate the problem areas. We have produced a list of the least efficient SQL queries issued by NELS and are addressing the methods needed to restructure the queries for more efficient operation.

Furthermore, we have used a variety of SQL scripts to track the usage of disk space. These statistics provide a scalable model of mass storage consumption for NELS based upon the database known as "christie." Using this model, we can project the storage required for

varying amounts of data and number of users. More work is needed to test the projections of this model.

We modified the distribution of objects in the christie tablespaces to allow the creation of all of the indexes. This substantially improved the query performance for that test data.

We have made a written recommendation (attached) of the target hardware for the ASV3 system. Note that system-level tuning cannot take place effectively until the target hardware is available.

2. *Investigate methods to increase the speed of NELS code and to provide additional improvement to the reset of the system. This includes modification of the NELS X-Windows code to interact with the Proc*C code more efficiently.*

The existing Pro*C code has two principal areas of inefficiency: 1) The code includes a number of badly phrased SQL statements; 2) the code does not reuse cursors.

We have modified a number of the most resource intensive and least efficient SQL commands. While the modified commands do execute much faster, more work is required to determine whether they still fall within the design requirements of NELS.

We have drafted several data models for NELS (attached). These reflect both the current data model and potentially improved data models. A change in the data model can achieve greater performance. At present we see no overwhelming reason to make major schema changes prior to the June delivery.

3. *Restructure existing code to interact with Oracle more effectively.*

We have modified some of the most frequently used SQL select commands to reuse cursors. This reduces the amount of time required to parse and execute these commands.

We obtained and assisted with the installation of the latest patch to Oracle 6.0.33, which corrects a number of Oracle bugs that may affect NELS. The new installation will be made operational soon by rebuilding databases on it and resetting environment variables to point to it.

We discovered a problem in the population of the history table. The date and time of each event are stored as a char(12) in the database, but are passed in as a long integer in C. The result is that leading zeroes are truncated, making the interpretation of the data and time problematic. We recommend either converting the field to an Oracle "date" data type or adding a redundant date data type field.

We met with Mark Rorvig to discuss the feasibility of modifications to the NELS schema and design that are desired in the next version of the system.

4. *Addition of error reporting code to help detect and remove bugs.*

We have added a standard error detection and notification facility to the NELS code. Using the error detection mechanisms, we have already detected execution problems that previously resulted in an unexplained blank screen.

This routine will display the error message, the SQL code that caused the message, and the Pro*C module and line number of the code causing the Oracle error. The error message

can be displayed both as an X-Window modal popup message and as text on the parent process's terminal.

5. *Recommend report writing tools to integrate with the ASV3 system.*

We have received a document identifying the reports desired by Mountain Net. We have installed SQL*Reportwriter and SQL*Forms V3 to utilize in meeting additional ASV3 requirements. Work in this area will accelerate in June.

NELS Initial Configuration—Estimates

Real Memory

To ensure all processes are in real memory (no paging):

Required by...	Meg. each	Number	Total Meg.
Sun Unix	16	1	16
Each X User	3	4	12
Each Oracle instance (2 recommended)	1	2	2
Each Oracle background process (4 per instance)	1	2x4	8
Oracle System Global Area (shared memory, per instance)	10	2	20
Each NELS user (est.)	3.5	11	39
Grand Total			97 meg.

Mass Storage

A critical optimization approach for Oracle database is to separate those data objects that are expected to be accessed simultaneously across different physical disk drives (spindles). In addition to placing tables and indexes on separate drives, individual Oracle tables and indexes can be "striped" across disk drives by the DBA to obtain benefits of parallel access. RAID (redundant arrays of inexpensive disks) configurations are strongly discouraged for Oracle databases.

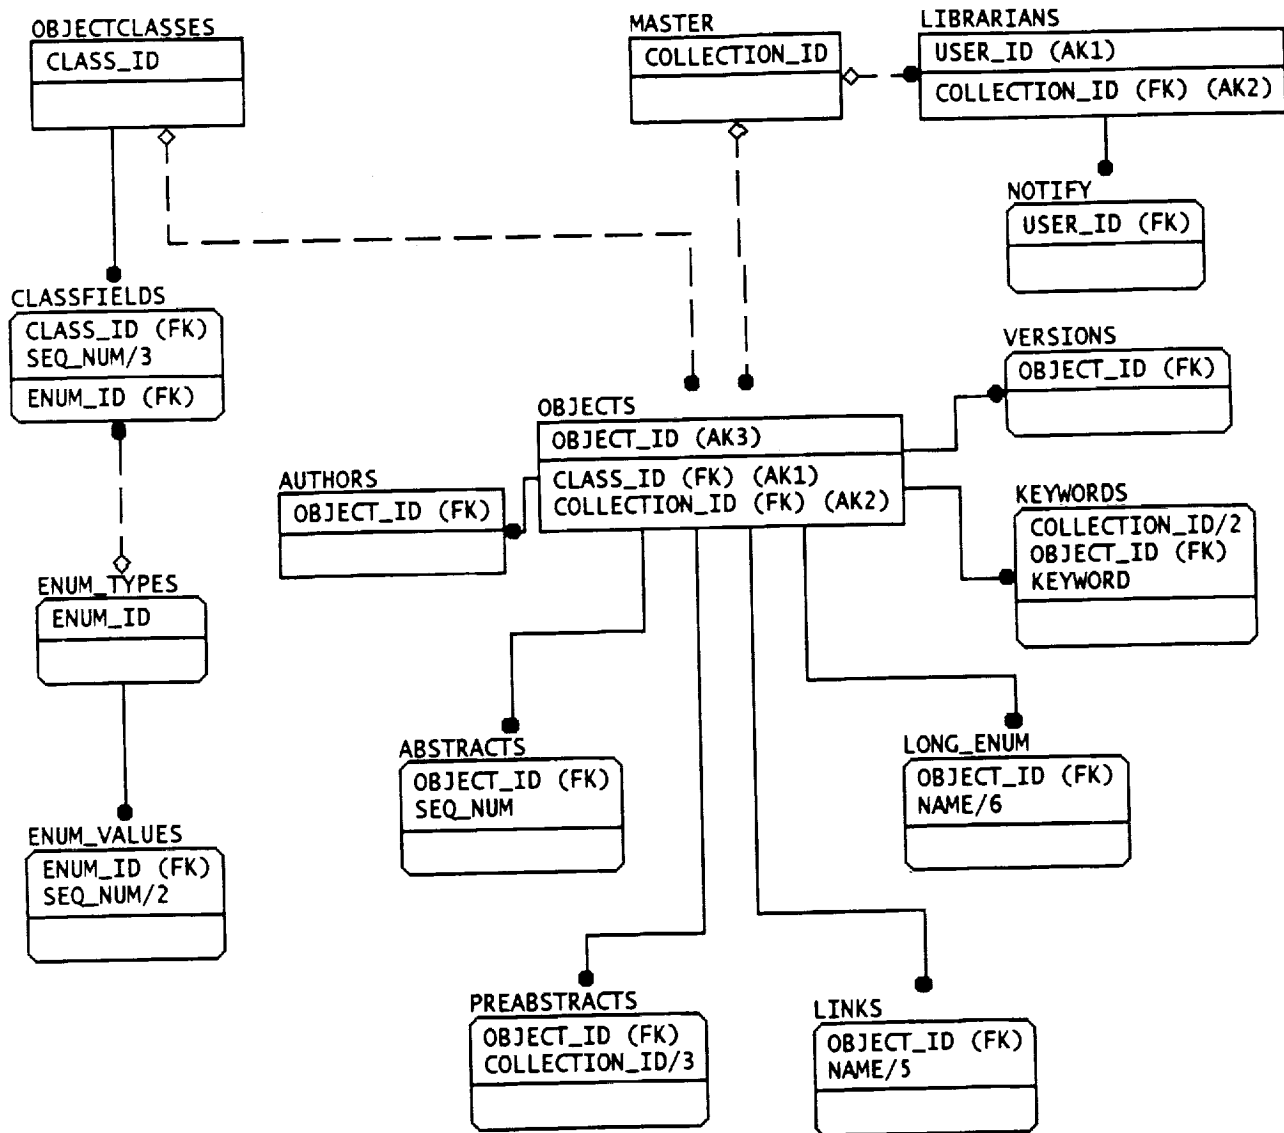
Unit	Contents	Size
1	Unix root partition; swap	500M
2	Oracle executable images; swap; library objects such as source code and gif images	1G
3	Oracle online redo log files—dedicated disk	500M*
4	NELS database files (Oracle tablespaces); library objects	500M
5	NELS database files; Oracle rollback segments	500M
6	NELS database files; Oracle temporary storage	500M
7	NELS database files; online backups; archives	1G

*Will require considerably less than size indicated. It is best if this drive is as fast as possible, to minimize waiting on writes to the database (all commits are journaled to a redo log file before control is returned to the requesting process).

CPU

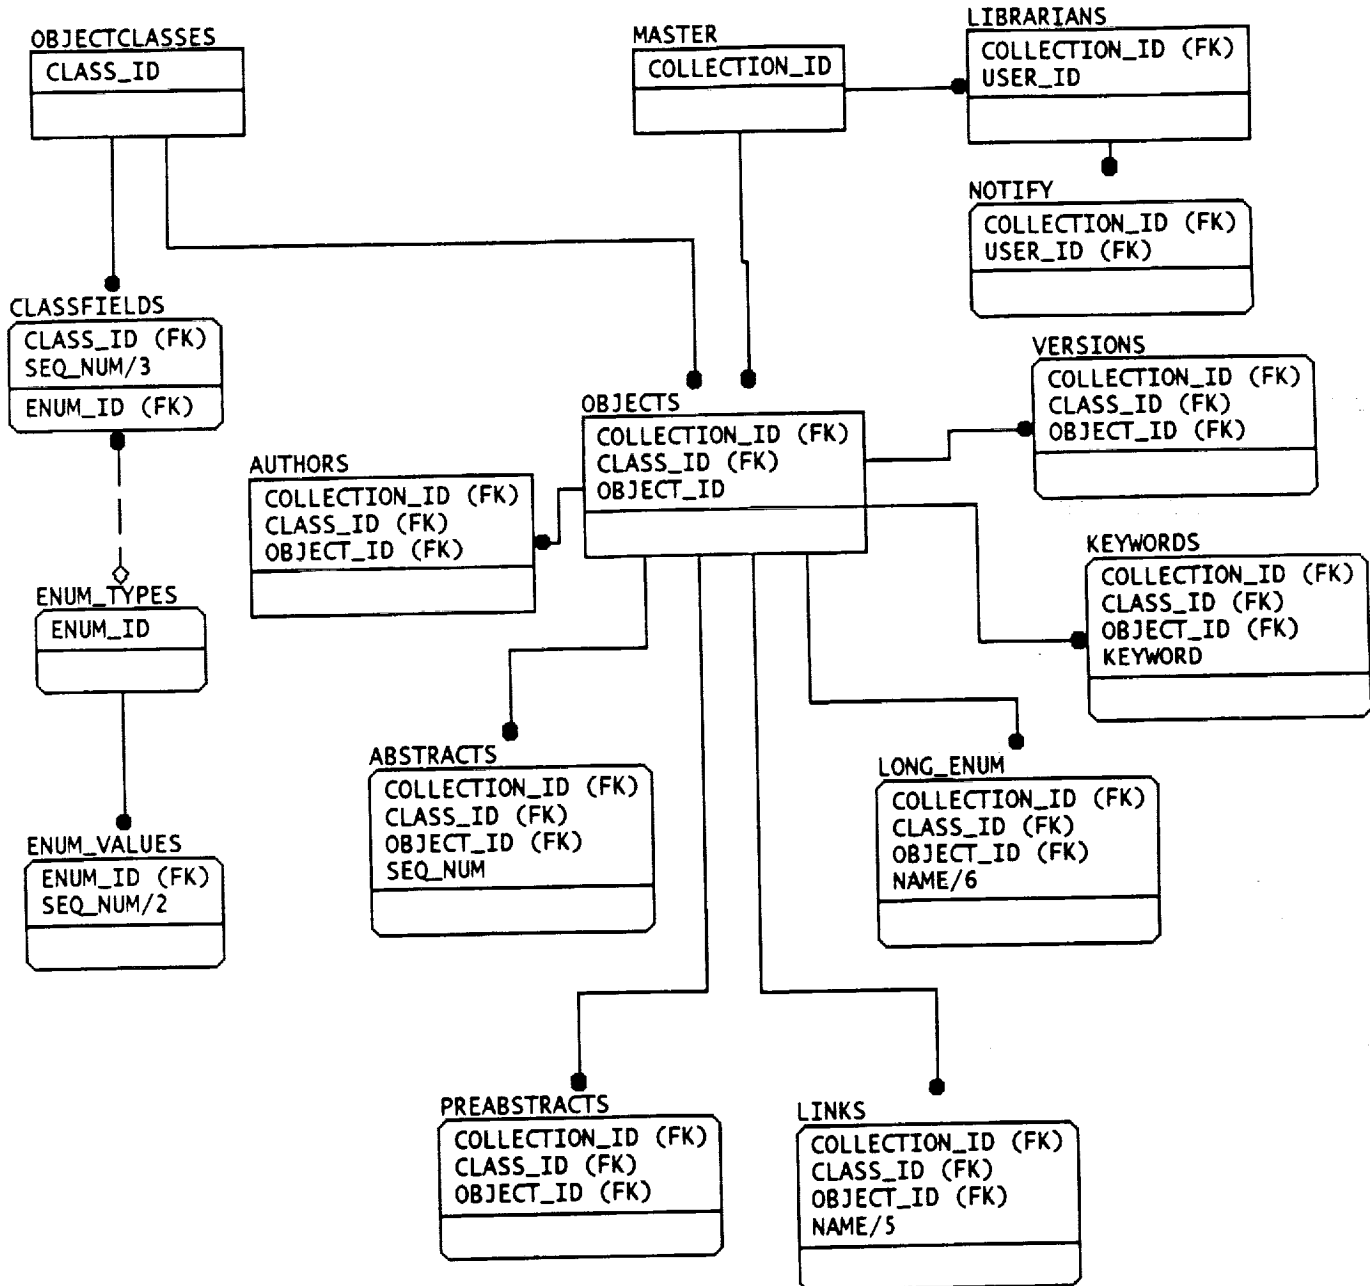
At least a two-processor machine is recommended. The Oracle DBMS can be configured to take advantage of multiple processors by setting the CPU_COUNT startup parameter.

NELS SCHEMA
MAY 28, 1992
EXISTING



NELS SCHEMA
MAY 28, 1992

ALTERNATIVE



NELS Optimization: June 92 Incremental Progress/Problem Report

Issues in Order of Importance

1. DataCraft has delivered versions of dblib and other C-language routines implementing many of the changes that we have recommended. We have also volunteered to implement the changes in the production version of the code, but because the production C-language code is out of our control, it is not known to us which of the changes will be implemented prior to delivery to Mountain Net. Within the bounds of the time remaining in the subcontract and extensions, we are available to perform the integration and testing of our delivered code.
2. Certain tasks, such as the installation of Oracle on the target hardware and distribution of physical database structures on target disk systems, could not be performed because the hardware was not yet available. DataCraft is willing to perform these tasks work via extension to this subcontract.
3. Although DataCraft has delivered tools to generate the high priority reports requested by Mountain Net, several of the reports require the collection of certain history data, not all of which was present in the data samples used in coding the reports. The code that produces history table entries was scheduled for modification by C.J. Mellebec/INET to collect desired historical data. In addition, Mountain Net has requested that a future version of NELS/ASV3 track historical data of group objects' members, as opposed to tracking only information about the group itself.
4. Access to the delivered SQL*Forms, SQL*Menu, and SQL*Reportwriter applications require access to an Oracle user id. The Nels application has so far been developed in the absence of Oracle-enforced security; that is, all privilege checking and enforcing is performed by the C code using Unix user id for authentication. It is our understanding that the Oracle user id and password are hardcoded somewhere on disk that is world-readable. Releasing the "mountainnet" Oracle user password compromises the security of the database, since all the structures are owned by this user. A decision is required from Mountain Net regarding the level of security desired in the system vs. the cost to implement.
5. Earlier estimates of real memory required by the X-windows version of NELS were probably low. There was some surplus in the amount of memory procured for the target hardware, and we are optimistic that memory constraints will not cause excessive performance problems. However, there is a small chance that additional memory may be required; testing of the target system will answer this question.
6. Access to workstations in the lab continues to be a problem. Often, upon arrival in the lab, every seat is filled, causing costly delays. Specifically, we need access to algol via VT-style terminal with known VT-style keyboard emulation. This problem could

Project Identification

Research Activity Number RB.07, Subcontract 110, Cooperative Agreement No. NCC 9-16. NASA Electronic Library System (NELS).

References

"NELS Optimization: May 92 Progress/Problem Report." Delivered by DataCraft, Inc. to the RICIS Document Control Department 29 May 92.

probably be solved by ensuring that every PC in the lab is networked and loaded with ftp Software's "rloginvt" emulator. We have addressed this problem by using our own equipment during these times.

7. Periodically, one or more of algol's critical file systems fills to capacity, stopping all work on the computer until there is intervention by a system manager. We recommend putting /tmp and /var on adequately sized file systems that are separate from root.
8. Providing a convenient method to edit abstracts via SQL*Forms requires the conversion of the abstract field to the Oracle LONG data type. This conversion is being performed by other developers and needs to be concluded soon.

Progress on Items Listed in Statement of Work

Each item from the statement of work is listed in italics, with a progress report below it.

1. *Examine NELS database, particularly in the memory, disk contention and CPU to discover bottlenecks.*

After delivering the initial hardware requirements, we performed additional measurement of memory and mass storage requirements for NELS. While the earlier mass storage recommendations were appropriate, the memory required to run the X-windows version of NELS can be considerably higher than our earlier estimate, depending on the type and number of queries issued. We did not measure the memory requirements of the ASCII LIB version. As mentioned above, the chance of this causing a performance problem is small but still exists.

DataCraft has discovered that the primary performance problems are caused by the phrasing of SQL SELECT statements that the NELS application generates. For example, there are numerous instances of forcing case-independent searches by applying the UPPER or LOWER function to a database field. Applying any function to a field in an SQL where-clause or order-by suppresses the use of a database index, thereby causing time-consuming full table scans.

We recommend storing all data that may be involved in a where- or order-by clause in fixed case, either upper or lower, and recoding the application to eliminate the case conversions. Karen Fleming/Mountain Net has confirmed the feasibility of doing so from the user perspective, with the exception of the object TITLE field, which must contain mixed case. Because of the requirement to order by title in many cases, we recommend storing a parallel UPPER_TITLE field in fixed case for use in order-by and where-clauses. Only the uppercase version of this field would be used in any indexes.

Supporting this change will be a need to record in the metadata information about which fields employ which data types. The easiest way to do this is split the existing CHARACTER data type defined for the CLASSFIELDS.KIND field into three distinct types: MIXED, ALL UPPER, and ALL LOWER.

Because this recommended change has been deferred, and because the NELS application would benefit from increased I/O throughput, we propose to "stripe" the largest tables and indexes (those associated with objects and preabstracts) across five of the seven disk drives on the target hardware, in order to achieve the benefits of parallel disk head movement. We are confident that this will give performance gains because of the relatively random distribution of access into these tables.

2. Investigate methods to increase the speed of NELS code and to provide additional improvement to the reset of the system. This includes modification of the NELS X-Windows code to interact with the Proc*C code more efficiently.

Using a copy of NELS separate from the production version, we have implemented and tested modifications to the Pro*C code which will yield performance improvements. We delivered a document to C.J. Mellebec/INET describing the changes we made and how to integrate the changes into the production code.

We have researched the issue of performing Control-C style break command for the X-Windows/Motif version of NELS. Unfortunately, it does not seem to be possible to support this feature, but we research on this issue continues at a low level.

We have also researched the issue of reducing memory consumption by using dynamic linking and shared code. While these features are supported by the current release of Sun's operating system, they are not yet supported by Oracle. This will likely change at some point in the future.

DataCraft also delivered code to add the ability for NELS to store ticks and ampersands in the database. Linked with this change is successful conversion of many dblib routines from dynamic SQL to static SQL.

At our suggestion, the action_date field in the history table has been converted to an Oracle DATE data type to facilitate searching and reporting.

As a result of meetings with Karen Fleming and C.J. Mellebec, we proposed several schemes that will make NELS easier to use. One is to store and display to the user the number of objects in each collection and all of its subcollections prior to the user executing any searches. Because of the roughly linear relationship between number of objects and search time, this would allow a user with only a moderate amount of experience to know in advance how quickly the result will appear. Drawing on this enhancement, it would be possible to have a pop-up warning when the user requests a query expected to exceed a certain threshold.

Another suggestion is to add a supplemental method of performing "like" searches. At present, these searches are executed over the entire word-space of the preabstracts, and can be quite time-consuming. An alternative way to do this would be to perform the search only on the keyword space; that is, a like-search would mean "show me all the objects with keywords like this object's keywords" instead of "show me all the objects with a preabstract like this object's preabstract." While it would not achieve an identical result, it may nevertheless produce a *useful* result, particularly if the user does not have a requirement to retrieve 100% of the potential matches.

Finally, we endorse an idea proposed by Mark Rorvig regarding an alternate protocol for interacting with the database. By first retrieving only a *count* of the number of hits achieved by a given request, instead of a large amount the metadata, it should be possible to reduce overall search time. The user could then rephrase his request to retrieve a manageable number of objects. In addition to memory savings, this approach could, if properly implemented, reduce the need for Oracle to build temporary segments for sorting, which is a resource-intense activity when performed on-disk.

3. Restructure existing code to interact with Oracle more effectively.

Queries generated by NELs sometimes include redundant subqueries and where clauses that slow processing. For example, consider the following statement issued by NELs:

```
select unique(object_name)
  from objects
 where (collection_id like 'S%'
        and class_id like '1%'
        and upper(object_name) = 'SHUTTLE')
 order by upper(object_name);
```

Since all collection IDs begin with S, all class ids begin with 1, and there is at most one row returned (making the order by irrelevant), this query would be better phrased as

```
select unique(object_name) from objects
 where upper(object_name) = 'SHUTTLE';
```

Similarly, the following query:

```
select unique(keyword)
  from keywords
 where (upper(keyword) = 'REQUIREMENT')
 and object_id =
 any (select unique(object_id)
       from objects
       where collection_id like 'S%'
         and class_id like '1%')
 order by upper(keyword)
```

is equivalent to the more efficient

```
select unique(keyword)
  from keywords
 where (upper(keyword) = 'REQUIREMENT');
```

While we did not have time to perform an exhaustive analysis of generated code, it is clear that in a number of cases the SQL is not as efficient as it could be.

4. Addition of error reporting code to help detect and remove bugs.

As mentioned in the last report, we have added a standard error detection and notification facility to the NELs code. The error message appears both as an X-Window modal popup message and as text on the parent process's terminal. During this contract period we have delivered to C.J. Mellebec a document which describes how one can integrate these routines into the production code.

5. Recommend report writing tools to integrate with the ASV3 system.

DataCraft made a selection of report writing tools by implementing reports of the highest priority as identified by Karen Fleming. To allow for ease of access and consistency of user interface into these reports, we constructed a SQL*Menu front end that invokes Bourne shell scripts for job control, which in turn call SQL*Plus and SQL*Reportwriter reports. In several cases, temporary tables are used to efficiently obtain the desired data. The reports implemented are the following:

- Object report
 - List of collections - alphabetic
 - List of collections - hierarchic
- Keywords
 - Objects accessed by collection - detail
 - Objects accessed by collection - summary
- Class definition

Objects created/modified/archived/deleted - detail
Most accessed objects
Least accessed objects

Note that the SQL*Menu application assumes the existence of certain environment variables. In addition to standard Oracle environment variables ORACLE_SID, ORACLE_HOME, ORAKITPATH, MENU5PATH, ORATERM, and LANGUAGE, the variable NELS_HOME must be set. Further, the shell scripts and SQL*Reportwriter files must exist in the \$NELS_HOME/bin directory, and the user's Unix PATH variable must include \$ORACLE_HOME/bin.

DataCraft has also delivered the SQL*Forms application requested by Mountain Net that will assist librarians in the cataloguing of objects. In addition to the conversion of the abstract field to a LONG data type mentioned previously, a number of other minor enhancements to the application are in progress.

APPENDIX: Location of Delivered Files

Path names of source code delivered on algol:

/home1/jfire/dblib	Modified Pro*C database access routines (all have been modified)
/home1/jfire/source	Several modified C language modules
/home1/jfire/includes	Modified include file
/home1/bpribyl/bin	Binaries for new SQL*Menu (.dmm), Forms (.frm), and ReportWriter (.rep) applications; Bourne shell scripts (.sh)
/home1/bpribyl/setup	SQL scripts to assist in the construction of the NELS tablespaces, tables, and indexes on the ASV3 hardware
/home1/bpribyl/source	ASCII export/dump files for SQL*Menu, Forms, and ReportWriter applications

NELS Optimization: July 92 Incremental Progress/Problem Report

Issues in Order of Importance

1. DataCraft is near the end of hours available in the subcontract. Although it is likely that NASA and/or Mountain Net will request enhancements that are beyond the limit of funding, it is not known what impact failure to perform such changes will have on the ASV3 system. Major tasks remaining include: *in situ* operational performance tuning; and installation of NELS on a second Sun system.
2. In delivering the next version of NELS to Mountain Net, NELS development personnel need to be cognizant of the need to preserve the existing NELS database, and deliver patch/update scripts that perform the minimum changes (for example, recompilations) required to install the update. More specifically: 1) The setup scripts `tables.sql` and `indexes.sql` have been obsoleted by the script `build.sql`; 2) Mountain Net will probably not be able to take the operational database offline for a rebuild.
3. Extensive work was required on the ASV3 production machine to compile and link NELS. It is not clear that this time consuming exercise was truly necessary since the production machine is binary compatible with the development machine. It would be much more efficient to distribute future releases in executable format whenever possible.
4. Because of the unavailability of the initial data set to obtain actual database size statistics, sizing in the initial database build on the production hardware were based on estimates. By making the estimates on the high side, we have attempted to mitigate the operational impact of using estimates instead of actual figures.
5. Although we striped the largest tables and indexes (those associated with objects and preabstracts) across five of the seven disk drives on the target hardware, resources did not permit ensuring that the initial database load actually distributed data across the striped extents. Without additional work we cannot ensure that the benefits of parallel disk head movement on expensive queries will be realized in the initial release of ASV3.
6. Our recommended change in the NELS code (and data) to remove case conversions on SQL SELECT statements has not been made. To repeat from our previous report, this change will have the singular largest beneficial impact on search performance.

Progress on Phase 1 Tasks

Phase 1 tasks were completed in June. See the 30 June 92 Progress/Problem Report for details.

Project Identification

Research Activity Number RB.07, Subcontract 110, Cooperative Agreement No. NCC 9-16. NASA Electronic Library System (NELS).

References

Previous deliveries of "NELS Optimization: Progress/Problem Report." Delivered by DataCraft, Inc. to the RICIS Document Control Department, dated 29 May 92 and 30 June 92.

Progress on Phase 2 Tasks

Each item from the statement of work is listed in italics, with a progress report below it.

1. Install the Oracle database management system and needed tools on ASV3 target hardware.

DataCraft installed Oracle RDBMS, SQL*Plus, SQL*Forms 3.0, SQL*Menu 5.0, Pro*C, SQL*Reportwriter, SQL*Net, and Pro*Ada on Sun hardware running SunOS 4.1.2 at Mountain Net's Cheat Lake, West Virginia facility on August 3-4, 1992. During the installation, a number of non-Oracle NELS issues were addressed with the assistance of Jim Folz/WVU. For example, DataCraft spent several hours debugging the compilation and linking of NELS executables. Because of the time dedicated to solving these unanticipated problems, the configuration of Oracle*Mail was left as an exercise for on-site DBA personnel.

2. Build NELS physical database structures and tune for performance on the target hardware.

DataCraft performed sizing estimates and built the tablespaces, tables, and indexes required by NELS on the target hardware. Doing so resulted in new version of scripts in the directory \$NELS_HOME/setup: **make_db.sql** (to build the database and tablespaces) and **build.sql** (to build the tables, indexes, and sequence).

Operational tuning must be deferred until there is a load on the system.

*3. Program needed modifications on delivered Pro*C software, including integration and testing as required.*

DataCraft supported the integration of our Pro*C code enhancements into the version 1.2 of NELS. In addition, we supported the development of code that eliminates the need to store the NELS password on disk. This code changed allows a NELS site to employ "OPSS" Oracle logins.

4. Perform Oracle system-level performance monitoring and tuning.

System-level performance monitoring and tuning must necessarily be deferred until there is a load on the system.

*5. Modify the ASV3 SQL*Menu front-end and supporting shell scripts as per user requirements.*

As proposed in the previous report, we developed a series of menu-driven scripts that allow a NELS site to add NELS users and grant each appropriate privileges. So doing increases the integrity of the NELS system by taking advantage of the Oracle kernel-enforced security. The developed scripts allow a NELS administrator to easily build public synonyms, add OPSS users, grant privileges to OPSS users, and revoke privileges from OPSS users.

6. Program additional metadata reporting capability as per user requirements.

DataCraft delivered an additional report on total and average duration of NELS user logon periods. In addition, several minor bug fixes on the metadata reports were performed during this period.

7. Program needed enhancements on delivered SQL*Forms applications.

DataCraft enhanced the SQL*Forms application as per user requirements. In addition, we provided technical assistance in the operation of the application during the user testing phase.

8. Documentation and reporting of these tasks.

This report constitutes said documentation and reporting.

NELS Optimization¹ : Final Report, Phase 2 Tasks

Issues in Order of Importance

New issues:

1. Configuration management of delivered code has begun to be a problem. DataCraft submitted new copies of NELSMENU code into an area on the ASV3 machine that held the NELS system. We had been mistakenly informed that the subdirectory was the correct location for delivered code. Several days passed before DataCraft was informed that the subdirectory in question was in fact not valid, and that the users had not been using the new versions.

We recommend clarifying the protocol for delivery of new NELSMENU code. For example, MountainNet should specify a single location to hold the latest tested versions. DataCraft would have write privilege in this area. Upon delivering a new version, DataCraft would send an E-mail to a designated configuration manager, whose job it would be to distribute and install the new code in the appropriate location(s).

Outstanding issues:

1. Our recommended change in the NELS code to remove case conversion functions on SQL SELECT statements has not, to our knowledge, been made. It is not known to us whether this change is scheduled for development.

Earlier Progress on Phase 2 Tasks

Together with the previous delivery², this report constitutes the final documentation of Phase 2 tasks.

Progress on Phase 2 Tasks

Each item from the statement of work is listed in italics, with a progress report below it.

1. *Install the Oracle database management system and needed tools on ASV3 target hardware.*

DataCraft provided post-installation support to Mountain Net, which included providing a sample script to perform Oracle database exports to disk for backup purposes.

2. *Build NELS physical database structures and tune for performance on the target hardware.*

No new database structures were implemented since the previous report. There is still an insufficient load on the system to perform operational tuning.

¹Research Activity Number RB.07, Subcontract 110, Cooperative Agreement No. NCC 9-16. NASA Electronic Library System (NELS).

²"NELS Optimization: July 92 Progress/Problem Report," Delivered 9 Aug. 92 to RICIS Document Control Department.

3. *Program needed modifications on delivered Pro*C software, including integration and testing as required.*

No modifications to delivered Pro*C software were performed since the previous report.

4. *Perform Oracle system-level performance monitoring and tuning.*

System-level performance monitoring and tuning is deferred until there is a load on the system.

5. *Modify the ASV3 SQL*Menu front-end and supporting shell scripts as per user requirements.*

No changes were required.

6. *Program additional metadata reporting capability as per user requirements.*

DataCraft added the capability to filter the most/least accessed reports by user id of the accessing user, or by "librarians" vs. "nonlibrarians." We also corrected a problem with subtotalling in the Object-Audit report.

7. *Program needed enhancements on delivered SQL*Forms applications.*

DataCraft programmed the following enhancements:

1. Added the capability to view and edit keywords associated with objects via the Object data-entry Form;
2. Improved field-to-field navigation within the form for ease of data entry;
3. Revised node_type validation and list-of-values to use enumerated type table;
4. Revised the Form to update the object count fields in the MASTER table when inserting or deleting objects, or when changing an object's node type;
5. Revised grantsp.sh to grant librarians update privilege on object count fields in MASTER;
6. Revised grantpub.sh to allow all users to insert, update, and delete from history, notify, and output_queue tables.

8. *Documentation and reporting of these tasks.*

This report constitutes said documentation and reporting.



NELS Optimization¹ : Interim Report, Phase 3

Issues in Order of Importance

New issues:

None

Outstanding issues from Phase 2:

1. (Second month) Our recommended change in the NELS code to remove case conversion functions on SQL SELECT statements has not, to our knowledge, been made. It is not known to us whether this change is scheduled for development.

Closed issues from Phase 2:

1. The configuration management problem was solved by agreeing with MountainNet personnel that the location for delivered code would be bpribyl's home directory on the primary NELS machine (hostname "a"), with an e-mail sent to the system administrator when a new version was delivered.

Progress on Phase 3 Tasks

Each item from the Phase 3 proposal is listed in italics, with a progress report below it.

1. *Install the Oracle database management system and tools on one additional Sun computer in West Virginia. Assist with NELS installation on this computer.*

DataCraft provided remote assistance to MountainNet during a first phase of NELS installation on a Sun workstation intended for demonstration purposes. The installation is not yet complete.

2. *Confirm understanding of NELS performance objectives for the production system. Set up tools for the collection of performance-related statistics on the production NELS computer under load. Analyze results of this instrumentation. Perform Oracle system-level tuning using these results.*

No progress this period.

3. *Implement CRs and DRs submitted by MountainNet on the NELSMENU system.*

3.1 *Add capability to edit KEYWORD field on the object-edit form*

DataCraft added a "detail block" to the object-edit form that provides a keyword edit capability on a per-object basis.

¹Research Activity Number RB.07, Subcontract 110, Cooperative Agreement No. NCC 9-16. NASA Electronic Library System (NELS).

3.2 *Enhance form navigation for speed of data entry*

DataCraft added navigational triggers that allow simple navigation through the object-edit form using the carriage return key instead of previous-block and next-block.

3.3 *Add capability in Most/Least accessed reports to select librarian vs. nonlibrarian*

The most/least accessed reports now provide the capability to filter by "librarian" or "nonlibrarian."

3.4 *Revise node_type validation and list-of-values to use enumerated type table*

The edit-object form now correctly performs validation and list-of-values on the node_type field.

3.5 *Revise the Object Form to update the object count fields in the MASTER table when inserting or deleting objects, or when changing an object's node type. Revise grantsp.sh to grant librarians update privilege on object count fields in MASTER.*

The edit-object form now updates the object count fields. Librarians are now granted field-level privileges on appropriate fields in the MASTER table via a pull-down menu.

3.6 *Revise grantpub.sh to allow all users to insert, update, and delete from history, notify, and output_queue tables.*

The public grant script now generates appropriate grants to public on the three tables named above.

3.7 *Correct a subtotalling problem in the object audit report.*

The object audit report now correctly subtotals on a per-collection basis.

3.8 *Eliminate ORA-1017 and ORA-1005 errors when executing Special-Add User*

This problem could not be reproduced, so no corrective action was taken.

3.9 *Revise object-edit form to populate an object archive table prior to deleting an object*

Object deletions via the object-edit form now make record of the object ID and name in the deleted_objects table.

3.10 *Assist MountainNet in establishing Oracle database backup scripts and procedures.*

DataCraft transmitted a sample database export script suitable for modification and inclusion in a nightly "cron" job.

4. *Continuing support of delivered Pro*C software (dblib routines).*

No work required on this activity during this period.

5. *Continuing support of the NELSMENU application.*

Additional work performed during this period:

1. Built a new object-audit report that includes the class id. User may select print destination, class ID, beginning and ending dates, and Unix user id on which to report.
2. In the object-edit form, added field-level validation of long_enum fields to ensure the abstractor has not exceeded line or column maxima as defined in the classfields table.
3. Corrected a problem wherein object extended attributes did not appear on the pop-up list of attributes if the class ID was longer than 4 characters.
4. Corrected a problem wherein object extended attributes of kind=8 (long_enum) were being stored in the objects table rather than the long_enum table.
5. Diagnosed and corrected a problem whose symptom was the apparent failure of NELS to store long_enum extended attributes. The cause was that the "abstract" field in the long_enum table was incorrectly built as a LONG data type instead of CHAR(240).
6. Modified the Class Definition report to include the "numcolumns" and "kind" fields from the classfields table.
7. Answered via telephone several questions raised by MountainNet related to data entry and reporting.

6. *Documentation and reporting of above tasks..*

This report constitutes a major portion of said documentation. In addition, a number of e-mailed messages to MountainNet personnel have kept them frequently apprised of the status of their requested enhancements.

NELS Optimization¹ : Final Report, Phase 3

Issues in Order of Importance

New issues:

1. A new version of Oracle, 6.0.36, is available on the Sun platform. This version includes bug fixes in Oracle*Mail which are required by Mountain Net. However, installing the new version will require regression testing of the NELS family of applications (that is, any NELS application that accesses the Oracle database) to ensure compatibility between NELS and 6.0.36. The amount of testing required is not expected to put an excessive drain on resources; however, it is impossible to predict whether the new Oracle version will require modifications to NELS code. DataCraft is available to install Oracle 6.0.36 on algol, and subsequently on the Adanet machine, as soon as project management authorizes and schedules needed personnel for application regression testing.

Outstanding issues from Phase 2:

1. (Third month) Our recommended change in the NELS code to remove case conversion functions on SQL SELECT statements has not, to our knowledge, been made. It is not known to us whether this change is scheduled for development.

Progress on Phase 3 Tasks

For brevity, progress that was documented in "NELS Optimization: Interim Report, Phase 3," has not been duplicated in this Final Report. Together with the interim report, the items below document Phase 3 in its entirety.

Each item from the Phase 3 proposal is listed in italics, with a progress report below it.

1. *Install the Oracle database management system and tools on one additional Sun computer in West Virginia. Assist with NELS installation on this computer.*

No additional work performed on this task during this period.

2. *Confirm understanding of NELS performance objectives for the production system. Set up tools for the collection of performance-related statistics on the production NELS computer under load. Analyze results of this instrumentation. Perform Oracle system-level tuning using these results.*

DataCraft is still awaiting the go-ahead from MountainNet to perform this task.

¹Research Activity Number RB.07, Subcontract 110, Cooperative Agreement No. NCC 9-16. NASA Electronic Library System (NELS).

3. *Implement CRs and DRs submitted by MountainNet on the NELSMENU system.*

3.1 *Add capability to edit KEYWORD field on the object-edit form*

Task completed in the previous reporting period.

3.2 *Enhance form navigation for speed of data entry*

Task completed in the previous reporting period.

3.3 *Add capability in Most/Least accessed reports to select librarian vs. nonlibrarian*

Task completed in the previous reporting period.

3.4 *Revise node_type validation and list-of-values to use enumerated type table*

Task completed in the previous reporting period.

3.5 *Revise the Object Form to update the object count fields in the MASTER table when inserting or deleting objects, or when changing an object's node type. Revise grantsp.sh to grant librarians update privilege on object count fields in MASTER.*

A bug was corrected to ensure that the total_objects field contains only the number of production objects, rather than all objects.

3.6 *Revise grantpub.sh to allow all users to insert, update, and delete from history, notify, and output_queue tables.*

Task completed in the previous reporting period.

3.7 *Correct a subtotaling problem in the object audit report.*

Task completed in the previous reporting period.

3.8 *Eliminate ORA-1017 and ORA-1005 errors when executing Special-Add User*

This problem has not recurred (or its recurrence was not reported to DataCraft) and was not investigated further.

3.9 *Revise object-edit form to populate an object archive table prior to deleting an object*

Task completed in the previous reporting period.

3.10 *Assist MountainNet in establishing Oracle database backup scripts and procedures.*

Task completed in the previous reporting period.

4. *Continuing support of delivered Pro*C software (dblib routines).*

No work required on this activity during this period.

5. *Continuing support of the NELSMENU application.*

Additional work performed during this period at the request of MountainNet:

1. Modified object-edit form to store new history action codes (26,27,28, and 29) as appropriate for the specified promotions and demotions of objects.
2. Modified object-edit form: Added validation to the node_type field to restrict promotions and demotions to (hard-coded) specified sequences.
3. Added the node label to the Audit report.
4. Performed a mass update of the keywords table to upper case. This required writing a short SQL script that eliminated keywords differing only in case.
5. Responded to a telephoned question from MountainNet regarding feasibility of upgrading Oracle to version 6.0.36 to correct an Oracle*Mail problem.
6. Corrected a bug in object-edit form which prevented preabstracts associated with objects from being deleted when the object was deleted.
7. Investigated source of problem wherein object names were not displayed in their entirety in asciilib.
8. Corrected a bug in object-edit form that gave deceptive results when using the duplicate-record function key in SQL*Forms, then navigating to the previous record and back to the "duplicated" record. It appeared that all attributes had been duplicated when in fact they were not.
9. Corrected a bug in object-edit form that forced an update of object records when retrieving a set of objects, then changing a subset of the retrieved objects. Previously, all objects retrieved in the set appeared to have been updated, rather than those for which some attribute(s) actually changed.
10. Corrected a bug in object-edit form which occurred when a new record was inserted, then immediately updated without requerying from the database. It appeared to update successfully but was not in fact committed to the database.

6. *Documentation and reporting of above tasks..*

This report constitutes a major portion of said documentation. In addition, a number of e-mailed messages to MountainNet personnel have kept them frequently apprised of the status of their requested enhancements.

NELS Optimization¹ : Interim Report, Phase 4

Issues in Order of Importance

New issues:

None

Outstanding issues from Phase 2:

1. (Fourth month) Our recommended change in the NELS code to remove case conversion functions on SQL SELECT statements has not, to our knowledge, been made. It is not known to us whether this change is scheduled for development.
2. (Second month) A new version of Oracle, 6.0.36, is available on the Sun platform. This version includes bug fixes in Oracle*Mail which are required by Mountain Net. However, installing the new version will require regression testing of the NELS family of applications (that is, any NELS application that accesses the Oracle database) to ensure compatibility between NELS and 6.0.36. The amount of testing required is not expected to put an excessive drain on resources; however, it is impossible to predict whether the new Oracle version will require modifications to NELS code. DataCraft is available to install Oracle 6.0.36 on algol, and subsequently on the Adanet machine, as soon as project management authorizes and schedules needed personnel for application regression testing.

Progress on Phase 4 Tasks

Each item from the Phase 4 proposal is listed in italics, with a progress report below it.

1. *Assist with installation of the Oracle database management system and tools on one additional Sun computer in West Virginia. Assist with NELS installation on this computer.*
Provided remote support to MountainNet for the installation of Oracle and NELS on a Sun IPX workstation used to demonstrate NELS at the TriAda conference.
2. *Confirm understanding of NELS performance objectives for the production system. Set up tools for the collection of performance-related statistics on the production NELS computer under load. Analyze results of this instrumentation. Perform Oracle system-level tuning using these results..*
No work performed on this task during this period.
3. *Continuing support of the NELSMENU application.*
The following tasks were performed at the request of MountainNet personnel:
 - 3.1. Ran a previously developed SQL script to accurately update the object counts for each collection in the "master" table

¹Research Activity Number RB.07, Subcontract 110, Cooperative Agreement No. NCC 9-16. NASA Electronic Library System (NELS).

- 3.2. Implemented a length limit of 2400 characters for object abstracts when they are created or edited via nelsmenu's object-edit form
- 3.3. Enhanced the keyword mass change form to include support for mass deletions
- 3.4. Developed a "category" SQL*Plus report that prints a list of objects by collection, title, and abstract

4. *Continuing support of delivered Pro*C software (dblib routines).*

No work performed on this task during this period.

4. *Support a new installation and administration of Oracle for NELS on a Sun computer at the University of Houston—Clear Lake.*

Answered questions regarding disk space allocation and other Oracle-related issues that arose during the installation.

5. *Documentation and reporting of above tasks..*

This report constitutes a major portion of said documentation. In addition, a number of e-mailed messages to MountainNet personnel have kept them frequently apprised of the status of their support requests.



NELS Optimization : Interim Report, Phase 4

Issues in Order of Importance

New issues:

None

Outstanding issues from Phase 2:

1. (Fifth month) Our recommended change in the NELS code to remove case conversion functions on SQL SELECT statements has not, to our knowledge, been made. It is not known to us whether this change is scheduled for development.
2. (Third month) A new version of Oracle, 6.0.36, is available on the Sun platform. This version includes bug fixes in Oracle*Mail which are required by Mountain Net. However, installing the new version will require regression testing of the NELS family of applications (that is, any NELS application that accesses the Oracle database) to ensure compatibility between NELS and 6.0.36. The amount of testing required is not expected to put an excessive drain on resources; however, it is impossible to predict whether the new Oracle version will require modifications to NELS code. DataCraft is available to install Oracle 6.0.36 on algol, and subsequently on the Adanet machine, as soon as project management authorizes and schedules needed personnel for application regression testing.

Progress on Phase 4 Tasks

Each item from the Phase 4 proposal is listed in italics, with a progress report below it.

1. *Assist with installation of the Oracle database management system and tools on one additional Sun computer in West Virginia. Assist with NELS installation on this computer.*

No work performed on this task during this period.

2. *Confirm understanding of NELS performance objectives for the production system. Set up tools for the collection of performance-related statistics on the production NELS computer under load. Analyze results of this instrumentation. Perform Oracle system-level tuning using these results..*

No work performed on this task during this period.

3. *Continuing support of the NELSMENU application.*

The following tasks were performed at the request of MountainNet personnel:

- 3.1. Modified the keyword mass change form so that it will update the timestamp on any object whose keywords change, and to create appropriate records in the history table. This change was designed to

NELS Optimization : Interim Report, Phase 4

provide the information needed by the batch job that updates the preabstracts table.

- 3.2 Modified the object audit (by collections) report to include added and modified objects that were subsequently deleted. Previously, any object added or modified but then deleted was not appearing on the reports.
 - 3.3 Rebuilt the history table index to include the object_id, collection_id, Action, and Acct_date.
 - 3.4 Modified the keywords report to include the number of uses of each keyword (by all objects).
4. *Continuing support of delivered Pro*C software (dblib routines).*
No work performed on this task during this period.
 4. *Support a new installation and administration of Oracle for NELS on a Sun computer at the University of Houston—Clear Lake.*
No work performed on this task during this period.
 5. *Documentation and reporting of above tasks..*
This report constitutes a major portion of said documentation. In addition, a number of e-mailed messages to MountainNet personnel have kept them frequently apprised of the status of their support requests.

NELS Optimization: Interim Report, Phase 4

Issues in Order of Importance

New issues:

None

Outstanding issues from Phase 2:

1. (Sixth month) Our recommended change in the NELS code to remove case conversion functions on SQL SELECT statements has not, to our knowledge, been made. It is not known to us whether this change is scheduled for development.
2. (Fourth month) A new version of Oracle, 6.0.36, is available on the Sun platform. This version includes bug fixes in Oracle*Mail which are required by Mountain Net. However, installing the new version will require regression testing of the NELS family of applications (that is, any NELS application that accesses the Oracle database) to ensure compatibility between NELS and 6.0.36. The amount of testing required is not expected to put an excessive drain on resources; however, it is impossible to predict whether the new Oracle version will require modifications to NELS code. DataCraft is available to install Oracle 6.0.36 on algal, and subsequently on the Adanet machine, as soon as project management authorizes and schedules needed personnel for application regression testing.

Progress on Phase 4 Tasks

Each item from the Phase 4 contract extension is listed in italics, with a progress report below it.

1. *Assist with installation of the Oracle database management system and tools on one additional Sun computer in West Virginia. Assist with NELS installation on this computer.*

No work performed on this task during this period.

2. *Confirm understanding of NELS performance objectives for the production system. Set up tools for the collection of performance-related statistics on the production NELS computer under load. Analyze results of this instrumentation. Perform Oracle system-level tuning using these results.*

No work performed on this task during this period.

3. *Continuing support of the NELSMENU application.*

The following tasks were performed at the request of MountainNet personnel:

- 3.1. Completed and transmitted an impact analysis for the changes required to NELSMENU that would result from the implementation of secure (proprietary) collections. For reference, the analysis is attached to this report.

NELS Optimization: Interim Report, Phase 4

- 3.2 Performed troubleshooting on a performance problem observed when modifying large sets of records via the Object-Edit SQL*Forms application.
4. *Continuing support of delivered Pro*C software (dblib routines).*
No work performed on this task during this period.
4. *Support a new installation and administration of Oracle for NELS on a Sun computer at the University of Houston—Clear Lake.*
No work performed on this task during this period.
5. *Documentation and reporting of above tasks:*
This report constitutes a major portion of said documentation. In addition, a number of e-mailed messages to MountainNet personnel have kept them frequently apprised of the status of their requests for support.

At the request of Glen Houston, wrote and submitted a short article about research activity RB.07 for inclusion in the second comprehensive RICIS research report. The time required to write the article was not billed to the contract.

Proprietary (Secure) Collection Functionality Estimated NELSMENU Impact Analysis

Ref: "Proprietary (Secure) Collection Functionality" fax received by Bill Pribyl from MountainNet 8 Dec 92.

Tables		
<i>Requested Change</i>	<i>Impact</i>	<i>Hours</i>
HISTORY Add: PROP_COL_FLAG CLASS_ID	1. Modify table build script 2. Alter online table in production database. May require table rebuild or tablespace extension to fit into tablespace. 3. Alter online table in "play" database instance. 4. Modify four queries each in Object Audit and Object Audit by Class reports	4
DELETED_OBJECTS Add: Collection_id Collection_name Class_id Classname	1. Modify table build script 2. Alter online table in production database. May require table rebuild or tablespace extension to fit into tablespace. 3. Alter online table in "play" database instance. 4. Modify four queries each in Object Audit and Object Audit by Class reports	4
PROP_COLL_ACCESS New table	1. Modify table build script 2. Build online table 3. Modify script that issues grants to issue appropriate privileges to librarians etc.	2
Reports		
Object existence	Add a displayed field that flags those objects which are in proprietary collections	1
Object Audit	Add a displayed field that indicates at the collection-break level which collections are proprietary	2
Object Audit by Class	Add a displayed field that flags those objects which are in proprietary collections	2
Collection, alphabetical	Add a displayed field that indicates which collections are proprietary	1
Collection, hierarchical	Add a displayed field that indicates which collections are proprietary	1
Objects Accessed	1. Modify "user" parameter to accept an option meaning "users with access to proprietary collections" 2. Add new "password" parameter to SQL*Menu, shell script, and report 3. Modify the report to display the input password, if any	5

Users, detail	<ol style="list-style-type: none"> 1. Modify "user" parameter to accept an option meaning "users with access to proprietary collections" 2. Add new "password" parameter to SQL*Menu, shell script, and report 3. Modify the report to display the input password, if any 	4
Collection password by password (new)	<ol style="list-style-type: none"> 1. Create new SQL*Plus report 2. Make available from menu 	2
User access to proprietary collections by password (new)	<ol style="list-style-type: none"> 1. Create new unparameterized SQL*Plus report 2. Make available from menu 	1
User access to proprietary collections by user id (new)	<ol style="list-style-type: none"> 1. Create new unparameterized SQL*Plus report 2. Make available from menu 	1
Forms		
OBJECT	<ol style="list-style-type: none"> 1. Modify transaction triggers to populate the HISTORY.Prop_col_flag and HISTORY.Class_id on insert, update, delete 2. Modify object.inp to populate new DELETED_OBJECTS fields (Collection_id, Collection_name, Class_id, Classname) on delete 3. Modify validation logic to enforce rule that objects in proprietary collections may only have certain node_types 	6
KEYWORD Add PROP_COL_FLAG, CLASS_ID to HISTORY table	<ol style="list-style-type: none"> 1. Modify transaction triggers to populate HISTORY.Prop_col_flag and HISTORY.Class_id when creating history records that result from keyword changes 	2
Total hours (est.): 38		

NELS Optimization: Final Report, Phase 4

Issues in Order of Importance

New issues:

None

Outstanding issues from earlier phases:

1. (Seventh month) Our recommended change in the NELS code to remove case conversion functions on SQL SELECT statements has not, to our knowledge, been made. It is not known to us whether this change is scheduled for development.
2. (Fifth month) A new version of Oracle, 6.0.36, is available on the Sun platform. This version includes bug fixes in Oracle*Mail which are required by Mountain Net. However, installing the new version will require regression testing of the NELS family of applications (that is, any NELS application that accesses the Oracle database) to ensure compatibility between NELS and 6.0.36. The amount of testing required is not expected to put an excessive drain on resources; however, it is impossible to predict whether the new Oracle version will require modifications to NELS code. DataCraft is available to install Oracle 6.0.36 on algol, and subsequently on the Adanet machine, as soon as project management authorizes and schedules needed personnel for application regression testing.

Progress on Phase 4 Tasks

Each item from the Phase 4 contract extension is listed in italics, with a progress report below it. For brevity, this Final report does not repeat documentation of interim progress that appeared in earlier reports.

1. *Assist with installation of the Oracle database management system and tools on one additional Sun computer in West Virginia. Assist with NELS installation on this computer.*

No work performed on this task during this period.

2. *Confirm understanding of NELS performance objectives for the production system. Set up tools for the collection of performance-related statistics on the production NELS computer under load. Analyze results of this instrumentation. Perform Oracle system-level tuning using these results.*

No work performed on this task during this period.

3. *Continuing support of the NELSMENU application.*

The following tasks were performed at the request of MountainNet personnel:

- 3.1. Enhanced Object-Audit reports to eliminate duplicate lines that were appearing after a user changed an object more than once on the same day.
- 3.2. Investigated what appeared to be an anomaly in the count of deleted objects in the Object-Audit reports.
- 3.3. Enhanced the object-Audit by Class report to include objects inserted, updated, or archived and then subsequently deleted (if user sets the class_id parameter to ALL).

NELS Optimization: Final Report, Phase 4

- 3.4 Fielded questions regarding reconfiguring the test database to be a demo database.
- 3.5 Documented several temporary tables that were introduced into the NELS schema last Fall to accommodate reporting requirements. E-mailed the documentation to Karen Fleming.
- 4. *Continuing support of delivered Pro*C software (dblib routines).*
No work performed on this task during this period.
- 5. *Support a new installation and administration of Oracle for NELS on a Sun computer at the University of Houston—Clear Lake.*
No work performed on this task during this period.
- 6. *Documentation and reporting of above tasks.*
This report constitutes a major portion of said documentation. In addition, a number of e-mailed messages to MountainNet personnel have kept them frequently apprised of the status of their requests for support.

NELS Optimization: Final Report, Last Phase

Issues in Order of Importance

New issues:

None

Outstanding issues from earlier phases:

1. (Second month) In addition to Oracle Corporation's release of version 6.0.36 (item 3 below), Oracle7 is now in production on the Sun platform. Oracle7 offers performance improvements and other kernel-level enhancements which merit a review with respect to NELS' ongoing requirements.
2. (Ninth month) Our recommended change in the NELS code to remove case conversion functions on SQL SELECT statements has not, to our knowledge, been made. It is not known to us whether this change is scheduled for development.
3. (Seventh month) A new version of Oracle, 6.0.36, is available on the Sun platform. This version includes bug fixes in Oracle*Mail which are required by Mountain Net. However, installing the new version will require regression testing of the NELS family of applications (that is, any NELS application that accesses the Oracle database) to ensure compatibility between NELS and 6.0.36. The amount of testing required is not expected to put an excessive drain on resources; however, it is impossible to predict whether the new Oracle version will require modifications to NELS code. DataCraft is available to install Oracle 6.0.36 on algol, and subsequently on the Adanet machine, as soon as project management authorizes and schedules needed personnel for application regression testing.

Progress on Phase 4 Tasks

Please note that, in the interest of saving costs, progress that has been documented in previous reports has not been duplicated here.

Each item from the Phase 4 contract extension is listed in italics, with a progress report below it.

1. *Assist with installation of the Oracle database management system and tools on one additional Sun computer in West Virginia. Assist with NELS installation on this computer.*

No work performed on this task during this period.

2. *Confirm understanding of NELS performance objectives for the production system. Set up tools for the collection of performance-related statistics on the production NELS computer under load. Analyze results of this instrumentation. Perform Oracle system-level tuning using these results.*

No work performed on this task during this period.

3. *Continuing support of the NELSMENU application.*

Contacted MountainNet personnel to determine whether they had any additional requirements DataCraft could satisfy before the termination of this contract. Because there was no response, DataCraft has assumed that all outstanding obligations on this Research Activity have been satisfied.

4. *Continuing support of delivered Pro*C software (dblib routines).*

No work performed on this task during this period.

NELS Optimization: Final Report, Last Phase

5. *Support a new installation and administration of Oracle for NELS on a Sun computer at the University of Houston—Clear Lake.*

No work performed on this task during this period.

6. *Documentation and reporting of above tasks.*

This report constitutes a major portion of said documentation.